

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	11	(database\$1 or (data adj base\$1) or db\$1) same select\$4 same group\$4 same backup\$1 same director\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/03/08 10:45
L2	0	L1 same (server\$1 same instance\$1 same connect\$4)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/03/08 10:46
L3	1	identify\$4 same violat\$4 same message\$1 same type\$1 same select\$4 same consist\$4 same check\$4 same (database\$1 or db\$1 or (data adj base\$1)) same log\$4 same access\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2005/03/08 10:50

Welcome to IEEE Xplore®

- Home
- What Can I Access?
- Log-out

Tables of Contents

- Journals & Magazines
- Conference Proceedings
- Standards

Search

- By Author
- Basic
- Advanced
- CrossRef

Member Services

- Join IEEE
- Establish IEEE Web Account
- Access the IEEE Member Digital Library

IEEE Enterprise

- Access the IEEE Enterprise File Cabinet

 Print Format

Your search matched **3** of **1134355** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance** in **Descending** order.

Refine This Search:

You may refine your search by editing the current search expression or entering a new one in the text box.

 Check to search within this result set

Results Key:

JNL = Journal or Magazine **CNF** = Conference **STD** = Standard**1 A structured graphical tool for analyzing boundary scan violations**

Cogswell, M.; Mardhani, S.; Melocco, K.; Arora, H.;
Test Conference, 2002. Proceedings. International , 7-10 Oct. 2002
Pages:755 - 762

[\[Abstract\]](#) [\[PDF Full-Text \(521 KB\)\]](#) **IEEE CNF****2 Dynamic polling scheme based on time variation of network management information values**

Yoshihara, K.; Sugiyama, K.; Horiuchi, H.; Obana, S.;
Integrated Network Management, 1999. Distributed Management for the Networked Millennium. Proceedings of the Sixth IFIP/IEEE International Symposium on , 24-28 May 1999
Pages:141 - 154

[\[Abstract\]](#) [\[PDF Full-Text \(596 KB\)\]](#) **IEEE CNF****3 Monitoring timing constraints in distributed real-time systems**

Raju, S.C.V.; Rajkumar, R.; Jahanian, F.;
Real-Time Systems Symposium, 1992 , 2-4 Dec. 1992
Pages:57 - 67

[\[Abstract\]](#) [\[PDF Full-Text \(936 KB\)\]](#) **IEEE CNF**

Welcome to IEEE Xplore®

- [Home](#)
- [What Can I Access?](#)
- [Log-out](#)

Tables of Contents

- [Journals & Magazines](#)
- [Conference Proceedings](#)
- [Standards](#)

Search

- [By Author](#)
- [Basic](#)
- [Advanced](#)
- [CrossRef](#)

Member Services

- [Join IEEE](#)
- [Establish IEEE Web Account](#)
- [Access the IEEE Member Digital Library](#)

IEEE Enterprise

- [Access the IEEE Enterprise File Cabinet](#)

 [Print Format](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
 The ACM Digital Library The Guide

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

violati and messag and check and access and log and backup and group and file

Found 62,765 of 151,219

Sort results by

 [Save results to a Binder](#)
[Try an Advanced Search](#)

Display results

 [Search Tips](#)
[Try this search in The ACM Guide](#)
 [Open results in a new window](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale

1 [Practical byzantine fault tolerance and proactive recovery](#)

Miguel Castro, Barbara Liskov

November 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 4Full text available: [pdf\(1.63 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Our growing reliance on online services accessible on the Internet demands highly available systems that provide correct service without interruptions. Software bugs, operator mistakes, and malicious attacks are a major cause of service interruptions and they can cause arbitrary behavior, that is, Byzantine faults. This article describes a new replication algorithm, BFT, that can be used to build highly available systems that tolerate Byzantine faults. BFT can be used in practice to implement re ...

Keywords: Byzantine fault tolerance, asynchronous systems, proactive recovery, state machine replication, state transfer

2 [Replication in the harp file system](#)

Barbara Liskov, Sanjay Ghemawat, Robert Gruber, Paul Johnson, Liuba Shrira

September 1991 **ACM SIGOPS Operating Systems Review, Proceedings of the thirteenth ACM symposium on Operating systems principles**, Volume 25 Issue 5Full text available: [pdf\(1.60 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes the design and implementation of the Harp file system. Harp is a replicated Unix file system accessible via the VFS interface. It provides highly available and reliable storage for files and guarantees that file operations are executed atomically in spite of concurrency and failures. It uses a novel variation of the primary copy replication technique that provides good performance because it allows us to trade disk accesses for network communication. Harp is intended to be u ...

3 [ARIES: a transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging](#)

C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz

March 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 1Full text available: [pdf\(5.23 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file systems and transaction-based operating systems. ARIES has been implemented, to varying

degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d ...

Keywords: buffer management, latching, locking, space management, write-ahead logging

4 **A coherent distributed file cache with directory write-behind**

Timothy Mann, Andrew Birrell, Andy Hisgen, Charles Jerian, Garret Swart
May 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 2

Full text available: [!\[\]\(74d4806277d7e73349d8e8c0897931e9_img.jpg\) pdf\(3.21 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Extensive caching is a key feature of the Echo distributed file system. Echo client machines maintain coherent caches of file and directory data and properties, with write-behind (delayed write-back) of all cached information. Echo specifies ordering constraints on this write-behind, enabling applications to store and maintain consistent data structures in the file system even when crashes or network faults prevent some writes from being completed. In this paper we describe ...

Keywords: coherence, file caching, write-behind

5 **Programming languages for distributed computing systems**

Henri E. Bal, Jennifer G. Steiner, Andrew S. Tanenbaum
September 1989 **ACM Computing Surveys (CSUR)**, Volume 21 Issue 3

Full text available: [!\[\]\(47734e4656765d20df4fdbd5b7aff048_img.jpg\) pdf\(6.50 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

When distributed systems first appeared, they were programmed in traditional sequential languages, usually with the addition of a few library procedures for sending and receiving messages. As distributed applications became more commonplace and more sophisticated, this ad hoc approach became less satisfactory. Researchers all over the world began designing new programming languages specifically for implementing distributed applications. These languages and their history, their underlying pr ...

6 **A survey of rollback-recovery protocols in message-passing systems**

E. N. (Mootaz) Elnozahy, Lorenzo Alvisi, Yi-Min Wang, David B. Johnson
September 2002 **ACM Computing Surveys (CSUR)**, Volume 34 Issue 3

Full text available: [!\[\]\(7bc43b319a082987e20f7bf78f4bab80_img.jpg\) pdf\(549.68 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This survey covers rollback-recovery techniques that do not require special language constructs. In the first part of the survey we classify rollback-recovery protocols into *checkpoint-based* and *log-based*. *Checkpoint-based* protocols rely solely on checkpointing for system state restoration. Checkpointing can be coordinated, uncoordinated, or communication-induced. *Log-based* protocols combine checkpointing with logging of nondeterministic events, encoded in tuples call ...

Keywords: message logging, rollback-recovery

7 **Integrating security in a large distributed system**

M. Satyanarayanan
August 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 3

Full text available: [!\[\]\(2088942ccfedc84a0a076c3fee3541aa_img.jpg\) pdf\(2.90 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Andrew is a distributed computing environment that is a synthesis of the personal computing and timesharing paradigms. When mature, it is expected to encompass over

5,000 workstations spanning the Carnegie Mellon University campus. This paper examines the security issues that arise in such an environment and describes the mechanisms that have been developed to address them. These mechanisms include the logical and physical separation of servers and clients, support for secure communication ...

8 Protection and the control of information sharing in multics 

Jerome H. Saltzer

July 1974 **Communications of the ACM**, Volume 17 Issue 7

Full text available:  pdf(1.75 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The design of mechanisms to control the sharing of information in the Multics system is described. Five design principles help provide insight into the tradeoffs among different possible designs. The key mechanisms described include access control lists, hierarchical control of access specifications, identification and authentication of users, and primary memory protection. The paper ends with a discussion of several known weaknesses in the current protection mechanism design.

Keywords: Multics, access control, authentication, computer utilities, descriptors, privacy, proprietary programs, protected subsystems, protection, security, time-sharing systems, virtual memory

9 A Survey of Techniques for Synchronization and Recovery in Decentralized Computer Systems 

Walter H. Kohler

June 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 2

Full text available:  pdf(3.33 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

10 Understanding fault-tolerant distributed systems 

Flavin Cristian

February 1991 **Communications of the ACM**, Volume 34 Issue 2

Full text available:  pdf(6.17 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

11 Columns: Risks to the public in computers and related systems 

Peter G. Neumann

January 2001 **ACM SIGSOFT Software Engineering Notes**, Volume 26 Issue 1

Full text available:  pdf(3.24 MB)

Additional Information: [full citation](#)

12 Recovery Techniques for Database Systems 

Joost S. M. Verhofstad

June 1978 **ACM Computing Surveys (CSUR)**, Volume 10 Issue 2

Full text available:  pdf(2.32 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

13 Recovery management in QuickSilver 

Rober Haskin, Yoni Malachi, Gregory Chan

February 1988 **ACM Transactions on Computer Systems (TOCS)**, Volume 6 Issue 1

Full text available:  pdf(2.21 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This paper describes QuickSilver, developed at the IBM Almaden Research Center, which

uses atomic transactions as a unified failure recovery mechanism for a client-server structured distributed system. Transactions allow failure atomicity for related activities at a single server or at a number of independent servers. Rather than bundling transaction management into a dedicated language or recoverable object manager, Quicksilver exposes the basic commit protocol and log rec ...

14 Distributed systems - programming and management: On remote procedure call

Patrícia Gomes Soares

November 1992 **Proceedings of the 1992 conference of the Centre for Advanced Studies on Collaborative research - Volume 2**

Full text available:  [pdf\(4.52 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The Remote Procedure Call (RPC) paradigm is reviewed. The concept is described, along with the backbone structure of the mechanisms that support it. An overview of works in supporting these mechanisms is discussed. Extensions to the paradigm that have been proposed to enlarge its suitability, are studied. The main contributions of this paper are a standard view and classification of RPC mechanisms according to different perspectives, and a snapshot of the paradigm in use today and of goals for t ...

15 Decentralized storage systems: Farsite: federated, available, and reliable storage for an incompletely trusted environment

Atul Adya, William J. Bolosky, Miguel Castro, Gerald Cermak, Ronnie Chaiken, John R.

Douceur, Jon Howell, Jacob R. Lorch, Marvin Theimer, Roger P. Wattenhofer

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Full text available:  [pdf\(1.87 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

Farsite is a secure, scalable file system that logically functions as a centralized file server but is physically distributed among a set of untrusted computers. Farsite provides file availability and reliability through randomized replicated storage; it ensures the secrecy of file contents with cryptographic techniques; it maintains the integrity of file and directory data with a Byzantine-fault-tolerant protocol; it is designed to be scalable by using a distributed hint mechanism and delegatio ...

16 The process group approach to reliable distributed computing

Kenneth P. Birman

December 1993 **Communications of the ACM**, Volume 36 Issue 12

Full text available:  [pdf\(6.00 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: fault-tolerant process groups, message ordering, multicast communication

17 A structural view of the Cedar programming environment

Daniel C. Swinehart, Polle T. Zellweger, Richard J. Beach, Robert B. Hagmann

August 1986 **ACM Transactions on Programming Languages and Systems (TOPLAS)**,

Volume 8 Issue 4

Full text available:  [pdf\(6.32 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents an overview of the Cedar programming environment, focusing on its overall structure—that is, the major components of Cedar and the way they are organized. Cedar supports the development of programs written in a single programming language, also called Cedar. Its primary purpose is to increase the productivity of programmers whose activities include experimental programming and the development of prototype software systems for a high-performance personal computer. T ...

18 Access control for large collections

H. M. Gladney

April 1997 **ACM Transactions on Information Systems (TOIS)**, Volume 15 Issue 2

Full text available: [pdf\(482.88 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Efforts to place vast information resources at the fingertips of each individual in large user populations must be balanced by commensurate attention to information protection. For distributed systems with less-structured tasks, more-diversified information, and a heterogeneous user set, the computing system must administer enterprise-chosen access control policies. One kind of resource is a digital library that emulates massive collections of paper and other physical media for clerical, en ...

Keywords: access control, digital library, document, electronic library, information security

19 The Recovery Manager of the System R Database Manager

Jim Gray, Paul McJones, Mike Blasgen, Bruce Lindsay, Raymond Lorie, Tom Price, Franco Putzolu, Irving Traiger

June 1981 **ACM Computing Surveys (CSUR)**, Volume 13 Issue 2

Full text available: [pdf\(1.75 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



20 Log files: an extended file service exploiting write-once storage

R. Finlayson, D. Cheriton

November 1987 **ACM SIGOPS Operating Systems Review , Proceedings of the eleventh ACM Symposium on Operating systems principles**, Volume 21 Issue 5

Full text available: [pdf\(1.07 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A log service provides efficient storage and retrieval of data that is written sequentially (append-only) and not subsequently modified. Application programs and subsystems use log services for recovery, to record security audit trails, and for performance monitoring. Ideally, a log service should accommodate very large, long-lived logs, and provide efficient retrieval and low space overhead. In this paper, we describe the design and implementation of the Clio log service. Clio pr ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
 The ACM Digital Library The Guide

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

violati and messag and check and access and log and backup and group and file Found 62,765 of 151,219

Sort results by

 [Save results to a Binder](#)
[Try an Advanced Search](#)

Display results

 [Search Tips](#)
[Try this search in The ACM Guide](#)
 [Open results in a new window](#)

Results 21 - 40 of 200

Result page: [previous](#)[1](#)[2](#)[3](#)[4](#)[5](#)[6](#)[7](#)[8](#)[9](#)[10](#)[next](#)

Best 200 shown

Relevance scale

21 [Fault tolerance under UNIX](#)

Anita Borg, Wolfgang Blau, Wolfgang Graetsch, Ferdinand Herrmann, Wolfgang Oberle
 January 1989 **ACM Transactions on Computer Systems (TOCS)**, Volume 7 Issue 1

Full text available: [pdf\(1.97 MB\)](#)
 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The initial design for a distributed, fault-tolerant version of UNIX based on three-way atomic message transmission was presented in an earlier paper [3]. The implementation effort then moved from Auragen Systems1 to Nixdorf Computer where it was completed. This paper describes the working system, now known as the TARGON/32. The original design left open questions in at least two areas: fault tolerance for server processes and recovery after a crash were brie ...

22 [Fault Tolerant Operating Systems](#)

Peter J. Denning
 December 1976 **ACM Computing Surveys (CSUR)**, Volume 8 Issue 4

Full text available: [pdf\(2.69 MB\)](#)
 Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

23 [The making of an unmonitored 24 hour access computer lab](#)

Sarah Baker
 November 1993 **Proceedings of the 21st annual ACM SIGUCCS conference on User services**

Full text available: [pdf\(769.15 KB\)](#)
 Additional Information: [full citation](#), [citations](#), [index terms](#)

24 [Principles of transaction-oriented database recovery](#)

Theo Haerder, Andreas Reuter
 December 1983 **ACM Computing Surveys (CSUR)**, Volume 15 Issue 4

Full text available: [pdf\(2.48 MB\)](#)
 Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

25 [Third Generation Computer Systems](#)

Peter J. Denning
 December 1971 **ACM Computing Surveys (CSUR)**, Volume 3 Issue 4

Full text available: [pdf\(3.52 MB\)](#)
 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The common features of third generation operating systems are surveyed from a general view, with emphasis on the common abstractions that constitute at least the basis for a "theory" of operating systems. Properties of specific systems are not discussed except where examples are useful. The technical aspects of issues and concepts are stressed, the nontechnical aspects mentioned only briefly. A perfunctory knowledge of third generation systems is presumed.

26 Understanding the limitations of causally and totally ordered communication

David R. Cheriton, Dale Skeen

December 1993 **ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles**, Volume 27 Issue 5

Full text available: [pdf\(1.71 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Causally and totally ordered communication support (CATOCS) has been proposed as important to provide as part of the basic building blocks for constructing reliable distributed systems. In this paper, we identify four major limitations to CATOCS, investigate the applicability of CATOCS to several classes of distributed applications in light of these limitations, and the potential impact of these facilities on communication scalability and robustness. From this investigation, we find limited meri ...

27 Continual repair for windows using the event log

James C. Reynolds, Lawrence A. Clough

October 2003 **Proceedings of the 2003 ACM workshop on Survivable and self-regenerative systems: in association with 10th ACM Conference on Computer and Communications Security**

Full text available: [pdf\(682.62 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

There is good reason to base intrusion detection on data from the host. Unfortunately, most operating systems do not provide all the data needed in readily available logs. Ironically, perhaps, Window NT and its successor, Windows 2000, provide much of the necessary data, at least for security events. We have developed a host-based intrusion detector for these platforms that meets the generally accepted criteria for a good Intrusion Detection System. Its architecture is sufficiently flexible t ...

Keywords: auditing, intrusion detection, intrusion response, survivability

28 The evolution of Coda

M. Satyanarayanan

May 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 2

Full text available: [pdf\(441.35 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Failure-resilient, scalable, and secure read-write access to shared information by mobile and static users over wireless and wired networks is a fundamental computing challenge. In this article, we describe how the Coda file system has evolved to meet this challenge through the development of mechanisms for server replication, disconnected operation, adaptive use of weak connectivity, isolation-only transactions, translucent caching, and opportunistic exploitation of hardware surrogates. For eac ...

Keywords: Adaptation, Linux, UNIX, Windows, caching, conflict resolution, continuous data access, data staging, disaster recovery, disconnected operation, failure, high availability, hoarding, intermittent networks, isolation-only transactions, low-bandwidth networks, mobile computing, optimistic replica control, server replication, translucent cache management, weakly connected operation

29 Verifying Security

Maureen Harris Chehelyl, Morrie Gasser, George A. Huff, Jonathan K. Millen

30 The Rio file cache: surviving operating system crashes

Peter M. Chen, Wee Teck Ng, Subhachandra Chandra, Christopher Aycock, Gurushankar Rajamani, David Lowell

September 1996 **Proceedings of the seventh international conference on Architectural support for programming languages and operating systems**, Volume 31 , 30 Issue 9 , 5

Full text available:  pdf(1.12 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

One of the fundamental limits to high-performance, high-reliability file systems is memory's vulnerability to system crashes. Because memory is viewed as unsafe, systems periodically write data back to disk. The extra disk traffic lowers performance, and the delay period before data is safe lowers reliability. The goal of the Rio (RAM I/O) file cache is to make ordinary main memory safe for persistent storage by enabling memory to survive operating system crashes. Reliable memory enables a syste ...

31 Checkpoint repair for out-of-order execution machines

W. W. Hwu, Y. N. Patt

June 1987 **Proceedings of the 14th annual international symposium on Computer architecture**

Full text available:  pdf(840.89 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Out-of-order execution and branch prediction are two mechanisms that can be used profitably in the design of Supercomputers to increase performance. Unfortunately this means there must be some kind of repair mechanism, since situations do occur that require the computing engine to repair to a known previous state. One way to handle this is by checkpoint repair. In this paper we derive several properties of checkpoint repair mechanisms. In addition, we provide algorithms for performing check ...

32 LIMITS-a system for UNIX resource administration

A. Bettison, F. Adcock, P. Chubb, A. Gollan, C. Maltby

August 1989 **Proceedings of the 1989 ACM/IEEE conference on Supercomputing**

Full text available:  pdf(1.03 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The UNIX operating system, despite its emergence as a standard for supercomputer systems, lacks effective support for multiuser resource administration. The design and implementation of a decentralised resource administration system uniformly realisable across the wide variety of UNIX dialects presents a number of problems. Among these problems are potential violations of the UNIX design philosophy, preservation of the user process environment and adherence to industry standards

33 Recovery of on-line data bases (Panel)

A. B. Tonik

January 1971 **Proceedings of the 1971 26th annual conference**

Full text available:  pdf(782.00 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

There has been much publicity lately about the difficulties associated with data processing systems. Customers complain that it is almost impossible for them to correct what they think are mistakes in bills sent to them by a data processing installation. Another aspect of data processing installations, which is just as important, is how to maintain files in an environment where errors can be generated by hardware or software failures. This brings up the subject of check point/rest ...

34 Cryptographic tools: The dual receiver cryptosystem and its applications

Theodore Diamant, Homin K. Lee, Angelos D. Keromytis, Moti Yung

October 2004 **Proceedings of the 11th ACM conference on Computer and communications security**

Full text available: [pdf\(329.14 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We put forth the notion of a dual receiver cryptosystem and implement it based on bilinear pairings over certain elliptic curve groups. The cryptosystem is simple and efficient yet powerful, as it solves two problems of practical importance whose solutions have proven to be elusive before:(1) A provably secure "combined" public-key cryptosystem (with a single secret key per user in space-limited environment) where the key is used for both decryption and signing and where encryption can be esc ...

Keywords: digital signature, elliptic curves, key escrow, pairing-based cryptography, public key, puzzles, useful secure computation

35 Towards a more effective way of teaching a cybersecurity basics course 

Rose Shumba

June 2004 **ACM SIGCSE Bulletin , Working group reports from ITiCSE on Innovation and technology in computer science education**, Volume 36 Issue 4

Full text available: [pdf\(262.95 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

The Cybersecurity Basics course is an interdisciplinary course for the Criminology, Management Information Systems and Computer Science students at IUP. The course introduces computer security by focusing on host security. This paper describes laboratory exercises developed as part of a project to augment and improve on the teaching of the Cybersecurity Basics course. Nine Linux-based laboratory exercises were developed. A poster paper, based on the developed laboratory exercises was presented a ...

Keywords: cybersecurity, exercises, security, tools

36 Information protection methods: Display-only file server: a solution against information theft due to insider attack 

Yang Yu, Tzi-cker Chiueh

October 2004 **Proceedings of the 4th ACM workshop on Digital rights management**

Full text available: [pdf\(311.80 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Insider attack is one of the most serious cybersecurity threats to corporate America. Among all insider threats, information theft is considered the most damaging in terms of potential financial loss. Moreover, it is also especially difficult to detect and prevent, because in many cases the attacker has the proper authority to access the stolen information. According to the 2003 CSI/FBI Computer Crime and Security Survey, theft of proprietary information was the single largest category of los ...

Keywords: access, digital rights management, information theft, insider attack

37 Rethinking the design of the Internet: the end-to-end arguments vs. the brave new world 

Marjory S. Blumenthal, David D. Clark

August 2001 **ACM Transactions on Internet Technology (TOIT)**, Volume 1 Issue 1

Full text available: [pdf\(176.33 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This article looks at the Internet and the changing set of requirements for the Internet as it becomes more commercial, more oriented toward the consumer, and used for a wider set of purposes. We discuss a set of principles that have guided the design of the Internet, called the end-to-end arguments, and we conclude that there is a risk that the range of new requirements now emerging could have the consequence of compromising the Internet's original design principles. Were ...

Keywords: ISP, Internet, end-to-end argument

38 Achieving robustness in distributed database systems

Derek L. Eager, Kenneth C. Sevcik

September 1983 **ACM Transactions on Database Systems (TODS)**, Volume 8 Issue 3

Full text available:  [pdf\(2.33 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The problem of concurrency control in distributed database systems in which site and communication link failures may occur is considered. The possible range of failures is not restricted; in particular, failures may induce an arbitrary network partitioning. It is desirable to attain a high "level of robustness" in such a system; that is, these failures should have only a small impact on system operation. A level of robustness termed maximal partial operability ...

Keywords: concurrency control, network partitioning, robustness, serializability

39 Practical uses of synchronized clocks in distributed systems

Barbara Liskov

July 1991 **Proceedings of the tenth annual ACM symposium on Principles of distributed computing**

Full text available:  [pdf\(997.91 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

40 Reliability Issues in Computing System Design

B. Randell, P. Lee, P. C. Treleaven

June 1978 **ACM Computing Surveys (CSUR)**, Volume 10 Issue 2

Full text available:  [pdf\(3.95 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Results 21 - 40 of 200

Result page: [previous](#) [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)